BOG 6: ModSim Session 1

ASCR Workshop on Extreme Heterogeneity in HPC

23-25 Jan 2018

Zoom: https://lbnl.zoom.us/j/7940966896

Dial(for higher quality, dial a number based on your current location):
US: +1 646 558 8656 or +1 669 900 6833 Meeting ID: 794 096 6896

International numbers available:

https://lbnl.zoom.us/zoomconference?m=2z_j6z4ZlHnm8sNiVinyaxJFYxglsrS7

BOG 6.1 - ModSim Contributors

Moderator(s): Andrew Chien, David Donofrio, Maya Gokhale, John Leidel (scribe)

BOGists:

- David Donofrio
- Adolfy Hoisie
- Christian Engelmann
- Andrew Chien
- John Leidel
- Jeanine Cook
- Jeremy Wilke
- Kevin Barker
- Lingda Li
- Zhiling Lan
- Anastasiia Butko
- Antonino Tumeo
- Si Hammond
- Andreas Gerstlauer
- John Shalf
- Dilip Vasudevan
- EJ Park
- Shinjae Yoo

BOGists cont:

- Candy Culhane
- Katie Schuman
- Maya Gokhale
- Sam Williams
- John Shalf
- Rob Roser
- Shirley Moore
- Travis Humble
- Rich Carlson
- Esam El-Araby
- Sudhakar Yalamanchili
- Nick Wright

ModSim Capability Targets for Extreme Heterogeneity

From the FSD: *Modeling and simulation (ModSim) of existing and proposed computer architectures and systems have been longstanding pillars of the ASCR portfolio. In supporting these research directions, ASCR recognized the need to "meet performance, energy-efficiency, and resilience requirements of systems and applications at all scales—from embedded to exascale—recognizing their broad impacts to the larger computational science community in a range of research areas, including those affecting national security and domain sciences."*

<u>Why are we here?</u>

Brainstorm and discuss what Modeling and Simulation capabilities will be needed in the ***<u>2025-2035 timeframe</u>*** *to make increasingly heterogeneous hardware technologies useful and productive for science applications. This session will focus on* ***node-level*** *issues.*

BOG 6.1 ModSim Current Status

- Broad set of simulators in the field today… many overlapping capabilities
- A few from the FSD:
- CPU
  - Gem5 - Full system, includes cache and on-chip communication. Highly detailed, but you pay in KIPS
  - ZSim - Faster x86 simulator, efficient concurrency modeling, scale to hundreds of cores.
- Memory
  - DRAMSim - Cycle accurate DDR Model
  - HMCSim - Cycle accurate HMC model
  - NVMain - non-volatile memory model
- NoC Models
  - OpenSoC - Cycle accurate NoC HW generator
  - Connect - FPGA NoC generator
  - Booksim - Cycle accurate SW model

(Maybe several viewgraphs)

BOG 6.1 ModSim Current Status (cont.)

- Frameworks / Libraries
  - SST - Mix and match components
  - SystemC - Event driven simulation library
- Emulation Systems
  - LiME - FPGA based memory / accelerator emulator
- … this list barely scratches the surface
- Integration, validation, LTS, etc. all issues

BOG 6.1 ModSim Current Status

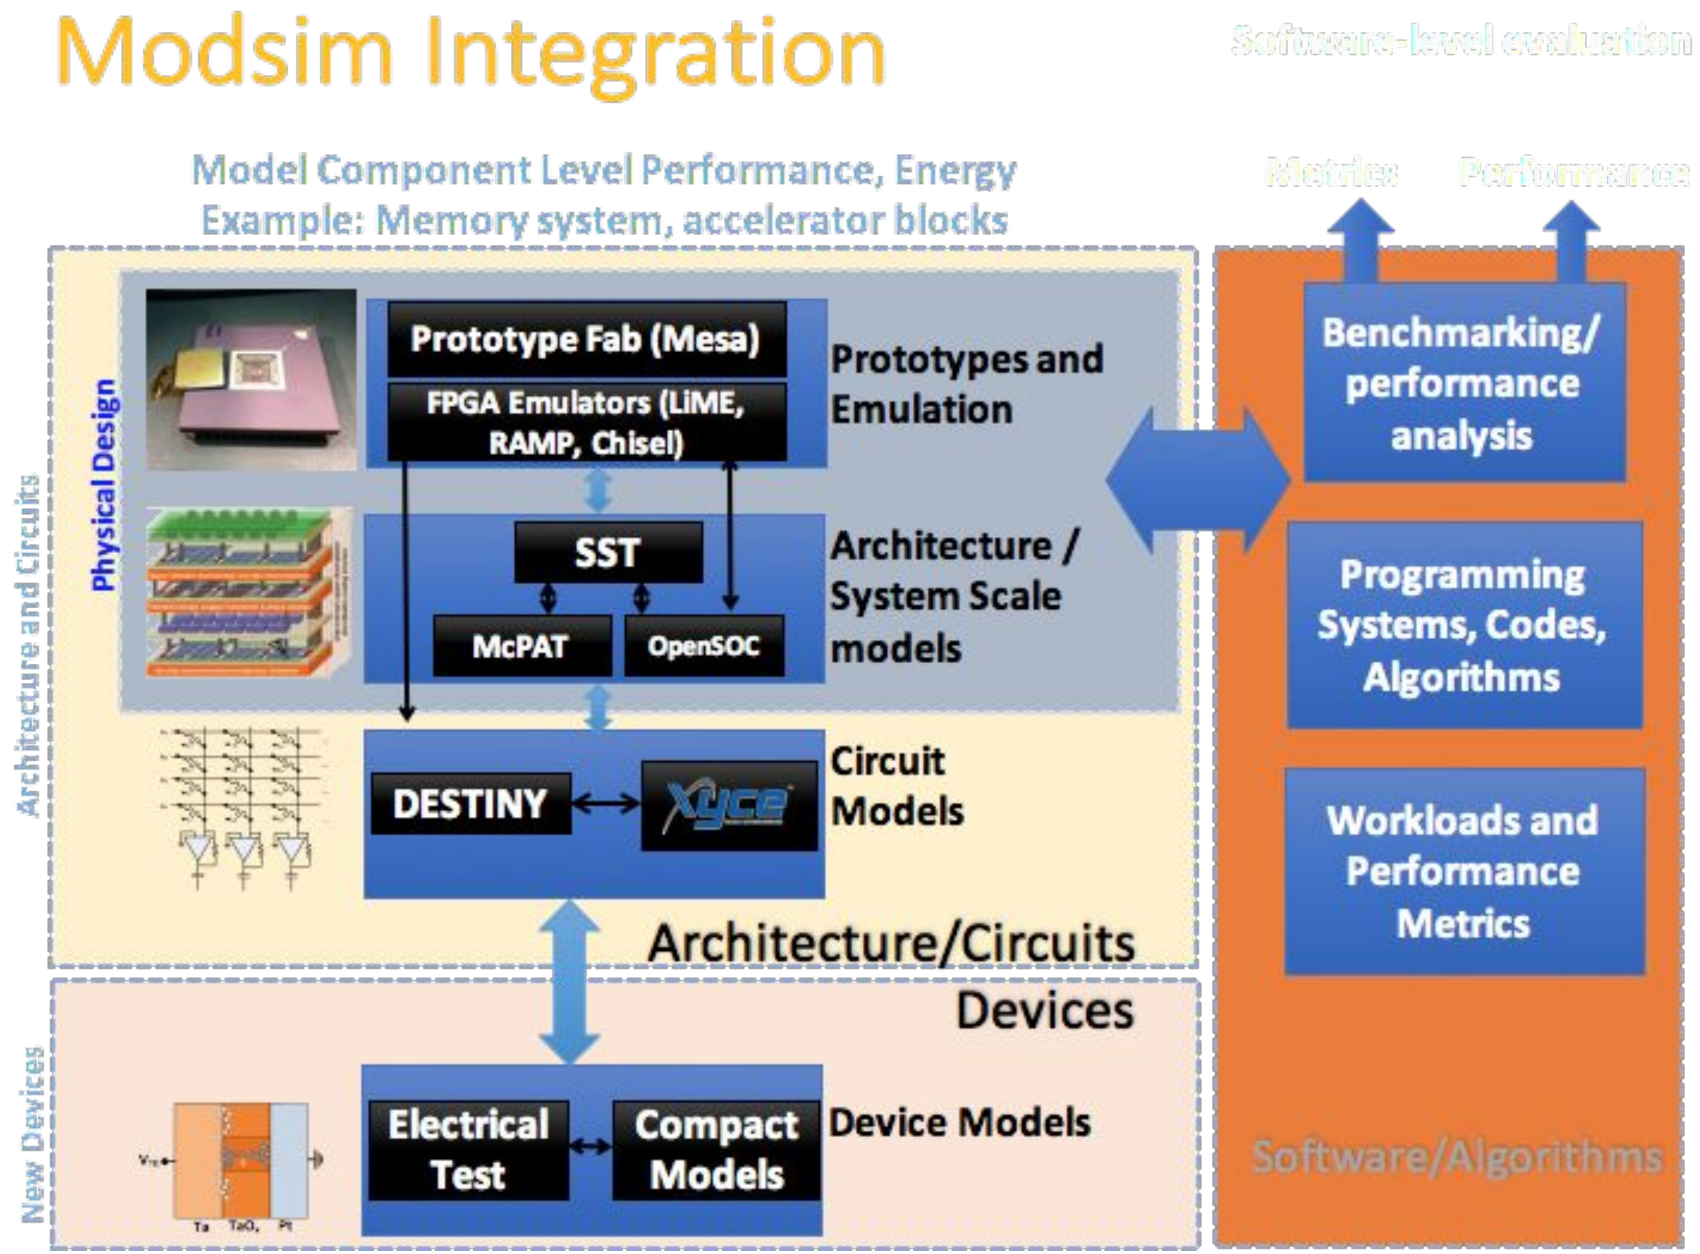

BOG 6.1 ModSim Themes from Whitepaper Submissions

- Performance Analysis Tools
    - On Node
    - Interconnect
- PostMoore Scaling
    - 3D Stacking
    - Device models
- Quantum Computing Models
- ML / Analytical modeling
    - Dark Silicon Management
    - Application Mapping
- Performance Models for Compilers
- FPGA Emulation Systems
    - Overlay
    - Memory / Accelerator Emulation
- Neuromorphic Processor models
- Smart network models
- Scalable Simulation Frameworks
    - Composability of models across simulation domains

Modeling and simulation - EH Requirements

EH requires a much broader and powerful capability:

- Quantitatively assess new architectural features
- Be capable of evaluating a wide range of granularities from microarchitecture to workflow level
- Be capable of integrating simulation components of novel heterogeneous accelerators such as quantum, neuromorphic, approximate, analog(ous)
- Properly assess characteristics of component interconnects
- Be scalable
- Characterize performance, power, thermal, reliability … attributes
- Need: Radically faster to explore a MUCH larger EH space

Node level modsim - EH Requirements

- Complex CPUs that incorporate wide variability in function units
  - Reduced precision, mixed precision, special purpose IP blocks
- Data visibility and sharing between closely coupled heterogeneous units
  - Accurately model complex protocols
  - Fast and scalable to enable comprehensive simulation (catch corner cases that show up non-deterministically and infrequently)
  - Model many different memory types
    - Potential combinatorial explosion of combinations
    - Compute in memory or near memory

ModSim EH BOG Capability Targets for 2030 (1/x)

Target 1: Increase accuracy while radically decreasing cost. Required to deal with explosion of EH architecture options. Goal: 10% error, at one-billionth modeling cost (performance cost)

Target 2: Full coverage of new technologies - Traditional electronic, Quantum, Analogous (whatever physical phenomena); and interoperation

Target 3: Modeling capabilities (standard, validated, range of speed/fidelity) that are integrated multi-scale simulation and modeling (1pS to 1 year; 1 operation to 1 Grand Challenge simulation)

Target 4: Cost of constructing and verifying the models and all derivations thereof

ModSim EH BOG Capability Targets for 2030 (2/x)

Target 5: Definition of new modeling & simulation methodologies that permit the prediction of future technology integration. Stochastic, ML-based

Target 6: Methodologies behind assembling models and integrating the analysis of these models against current and future EH technologies

Target 7: Integrated modeling of multiphysics of technologies alongside traditional modeling targets (et al. RAS, programmability, facilities)

Target 8: Validation of modeling techniques and core architecture/workload models for EH technologies in target time frame. Multifaceted validation across architectures, applications and data

ModSim EH BOG Capability Targets for 2030 (3/x)

Target 9: Dynamic modeling; Assembling models on the fly using information that we gather dynamically. Model introspection

Target 10: Deterministically building multiscale models that maintain a certain degree of performance and accuracy for the desired targets

Target 11: Utilizing emerging technologies as the basis for future systems modeling and simulation. Utilizing technologies such as quantum computing to model future systems; Emerging SPD’s (eg ML systems) to model future systems

Target 12: Modeling resource and management and scheduling as associated with node/system modeling (et al. architecture, workload modeling)

ModSim EH BOG Capability Targets for 2030 (4/x)

Target 13:

EH Node MODSIM Key Challenges

Need accuracy with radical decrease in cost (“scientific” style models)

- Not detailed (speed, effort), but accurate => deal with explosion of options
- => Goal: 10% error, at $1/10^9$ cost; AI/ML as key strategy
- Example: Accurate application-architecture potential performance estimation across EH space of dozens of dimensions
  - Selection amongst dozens of architectures (which are worth porting, even if the porting is automated)
  - Selection across dozens of attributes in one architecture (which features are worth exploiting, or just elide them)

More EH Node MODSIM Key Challenges

- Models (standard, validated, range of speed/fidelity) for new technologies (device, physical phenomena)
  - Traditional electronic, Quantum, Analogous (whatever physical phenomena)
  - Fully interoperable with CMOS-Digital models
  - => more than performance (per Colwell - power, communications, side channels, function, leakage)
- Modeling capabilities (standard, validated, range of speed/fidelity) that are integrated multi-scale simulation and modeling
  - Hardware and software phenomena -- Picoseconds (switch) to milliseconds (nvm) to seconds (heat) to months (wearout)
  - Application and multi-program behavior - nanoseconds (operations) to microseconds (thread and chip scale) to milliseconds (user) seconds () to minutes ()

More Node MODSIM Key Challenges

Facilities and Operating Systems

- Lightweight, accurate application-architecture performance estimation, and interference for multi-program workload
  - allocate and schedule applications across the 12 different accelerators (computational resource) types within the Yottascale machine
  - co-schedule applications on partitions, memory, nodes/accelerators, IO
  - => more than performance (per colwell - power, communications, side channels, function, leakage)

ModSim Key Challenge / Capability Directions - Discussion (1/x)

- **Topic 1**: Modeling future applications for target architectures that do not yet exist (and possibly utilize programming models that are in flux and/or not yet developed). EG, traditional trace-based models will not be sufficient for all future modeling methodologies
    - *Cross-Cutting Topic 1*: Interfacing with programming models group
    - *Cross-Cutting Topic 2*: Interfacing with performance portability group
    - *Cross-Cutting Topic 2:* Interface with operating system & system software group
- **Topic 2**: Ensuring the confidence (et al. accuracy) of the modeling and simulation methodologies for technologies that may utilize programming models and workflows that are in flux and/or not yet developed by prove to be difficult. How do we *trust* the models? Incorporating the physical characteristics of future technologies may inherently change the metrics of accuracy and feasibility.
- **Topic 3**: The engineering challenges around constructing multi scale models (especially when emerging hardware technologies are utilized for modsim) can be especially difficult. How do we ensure stable and reproducible modeling and simulation environments. While this may not result in a direct research direction, it needs to be considered as a rigorous engineering effort.

ModSim Key Challenge / Capability Directions - Discussion (2/x)

- **Topic 4:** Do we need reduced/abstract models as targets for compilers, runtime libraries, numerical libraries and system software? How do we accurately represent reduced function models sufficient to build the required infrastructure around the device? This may include multiple architectural targets with disparate ISAs and uArchs.
  - *Cross-Cutting Topic 1:* Interfacing with programming models group
- **Topic 5**: The ability to rapidly create **CORRECT** models. Model creation agility for functional models, analytical models, parallel models (in a compositional fashion). “Generation of models at all scales.” The need for non-determinism where appropriate is required. An accepted and expected degree of non-determinism in specific models (and multi scale simulations) will need to be dealt with appropriately and understand as it relates to accuracy of results.
  - *Cross-Cutting Topic 1*: Interfacing with programming models group
  - *Cross-Cutting Topic 2*: Interfacing with performance portability group
- **Topic 6**: Modeling/simulation of execution and programming environment for post-binary architectures
  - *Cross-Cutting Topic 1*: Interfacing with programming models group

ModSim Possible Research Directions Summary (1/x)

PRD X.1 - Need for research that covers the necessary methodologies and techniques encompassing all new emerging architectural techniques, physical phenomenon, devices, power/performance, reliability and facilities (see reference to BOG6.Session2). Cognisance of interoperability (standards-based).

PRD X.2 - Need for cohesive integration of modeling & simulation infrastructure with programming models/environment in multiple dimensions

PRD X.3 - Given the broad use of modsim for programming environments, architecture exploration and facilities planning, the performance (both in time to execute and time to create models) is paramount. *Reduce the effective cost of modsim*

ModSim Possible Research Directions Summary (2/x)

PRD X.4 - Need for research associated with applications/runtime modeling for future EH platforms and architectures. How do we model the future application spaces in a static and dynamic manner?

PRD X.5 - Relational modsim for integrating application/architecture modeling alongside workload/facility scheduling. How do we model sharing, provisioning, interaction, conflict between functional systems?

PRD X.6 - Research associated with modsim accuracy, reliability and feasibility. Confidently associating research results and defining metrics alongside models as the basis for feasible systems/workloads. Better integration of machine learning to improve modsim performance and lower cost. <ML is cross cutting across all PRDs>

ModSim Possible Research Directions Summary (2/x)

PRD X.7 - Emergent behavior

PRD 6.1 : Methodology and Tools for Accurate Modeling of Extreme Heterogeneity (Adolfy Hoisie, Sudhakar Yalamanchili, Zhiling Lan)

The current state of the art in modeling methodologies was developed in response to the architectures and application structures typically found in today's HPC landscape. However, as we move toward extreme heterogeneity and beyond Moore's Law, many of the assumptions built into these methods will no longer hold true. New models and modeling methodologies will need to be developed or current ones adapted to an increasingly diverse and complex landscape.The space under investigation is multi-dimensional: technologies from new accelerators, memories and storage, networks concepts and beyond Moore's law technologies, the triad performance/power/reliability, applications and algorithms designed for heterogeneity, and static and dynamic modeling.

- Research challenges:
  - Scope: Coverage of relevant new technologies, e.g., device, quantum, neuromorphic, analogous computing, etc.
  - Integrated Modeling: Development of methodologies and tools for integrated modsim of power/performance/reliability/thermals
  - Application Modeling: Ability to capture and model realistic application workloads that use very diverse compute and storage technologies
  - Multiscale: Modsim methodologies adapted for different scales of EH computing - from microseconds (power) through workflows (days)
  - Accuracy Tradeoffs: Modsim methodologies for EH at different levels of accuracy.
  - Dynamic Modeling: Methodologies for dynamic modeling: assembling models on the fly for runtime/OS/compiler introspection
  - Interoperability: Supporting interoperability via composition of workload, architecture, performance, and physical (e.g., power) models
  - Cost reduction: For similar or better accuracy and level of details for the integrated simulations through new methodologies, reduction of model generation and execution time, automation, and simplification of validation.
- Potential research approaches and research directions: Numerous challenges given a multidimensional space. Investigation of new stochastic methods including ML. Methodologies and tools to be researched and encapsulated in tools for modsim of new architectures.
- How and when will success impact technology? Success will be gradual as new methods and tools become available.
- Why make this a priority research direction? EH exhibits a large number of degrees of freedom in the way in which mapping workloads to system resources is done for optimality. Modsim is the key to achieving productive computing by many metrics. Design space exploration. Ability to optimize execution of apps on systems dynamically.
- Metrics:
  - Cost reduction, improved accuracy, ease of generation.

PRD 6.3 : Reduce the Effective Cost of modsim (1,000x?) (Andrew Chien, Shirley Moore) (1 of 2)

EH hardware complexity means that modeling and simulation will be an essential tool to manage complex choices - programming environments (compilation, runtime), resource management, facilities planning, architecture evaluation, etc. The cost (both time to create and execute models) is paramount, creating the ability to use good models for effective design and use of EH systems.

Research Challenges

- Efficient modeling - what needs to modeled in detail and what can be approximated (e.g., sampled, ignored, fast forwarded), how to approximate, how to combine different approaches
- Low-cost model development, cost cannot scale with EH applications X EH hardware
- Low-cost model validation

Research Approaches and Directions

- Partially-automated approaches for model building, including intelligent search and machine learning
- Partially-automated tool approaches for application model construction (e.g. skeletonization, sampling)
- DeepNeuralNetworks/ANN (accelerated) and other approximation techniques for low-cost model realization
- ML - observation, training, automated model extraction based on prior simulations and extrapolation
- Hardware acceleration of ML models, parallelism, Intelligent sampling based approaches

PRD 6.3 : Reduce the Effective Cost of modsim (1,000x?) (Andrew Chien, Shirley Moore) (2 of 2)

Metrics for success

- Time (human and machine) to develop node *and* system models of different complexity
- Time to develop application models
  - Emphasis on composability as critical to productivity
- Time to run models and obtain meaningful results

How and when will success impact technology?

- Embedded in programming environment tools, platform software (thru vendors), and used by facilities in planning, platform evaluation
- X-cut to Programming Environments, Scheduling and Resource Management

PRD 6.6 : Validation of Models of Extreme Heterogeneity (Sudhakar Yalamanchili)

We are faced with the need to model future systems, programming models, and application workloads that are i) new, ii) whose development is in a state of flux, iii) have yet to be developed, or iv) may emerge tomorrow. In such a dynamic environment, there is a need for the development of MODSIM methodologies and techniques for constructing correct models. The approaches should be designed to be agile and rapidly retargeted to accommodate the unanticipated (but likely to occur) emergence new technologies, systems, and workloads.

- Research challenges:
    - Metrics: Definition of metrics for driving modsim development. For example, how do we capture impact of technologies still under development?
    - Machine Learning: Better integration of ML techniques, for example to bridge modeling gaps via synthesis of performance models based on technology or workload projections
    - Accuracy: Methodologies and tools to ensuring trust and confidence in the MODSIM especially with regard to technologies (hardware and software) that are still under development
    - Validation Methodologies: Validation of modeling techniques and core architecture/workload models for EH technologies in target time frame. Multifaceted validation across architectures, applications and data
    - Emergent behavior: how to validate combined behavior of different components
- Potential research approaches and research directions: Targeted prototyping to assist in validation, borrow some techniques from model correlation techniques, inspiration from existing verification/validation techniques, and model reduction techniques
- How and when will success impact technology? The first set of validated models, perhaps applied to ECP machines as test vehicle
- Why make this a priority research direction? We need a high degree of confidence in the models that are developed and used to influence the national HPC agenda

PRD 6.1b??? : Impact for 6.1 is cohesive integration of modeling & simulation infrastructure with programming models/environment (Jeremiah Wilke, David Donofrio)

Programming models geared towards extreme heterogeneity will emphasize performance portability across many target architectures. The compilers and runtimes supporting each programming environment will have to decide between multiple target architectures while also tuning multiple parameters within each architecture. ModSim should help fill a capability gap by providing cost models. These cost models should aid 1) app developers using a programming environment and 2) compiler and runtime developers building the environment navigate the multi-dimensional EH space.

- Research challenges
  - *Accuracy tradeoffs:* Models with sufficient accuracy to inform “good” decisions, but efficient enough for compilers/runtimes
  - *New hardware technologies*: How to combine ModSim techniques for non-Von Neumann (neuromorphic, quantum) with techniques for “more Moore” architectures
  - *Interoperability*: Tools developed must be composable with programming environments
- Research approaches and research directions
  - ModSim helps environments: Online workload models integrated with on-node scheduling
  - ModSim helps environments: Machine learning, synthesis models, or mathematical methods for selecting compiler/runtime path from high-dimensional space
  - Environments help ModSim: Flexible workload models for ModSim tools to understand application performance on speculative technologies
- How and when will success impact technology?
  - ModSim will have a gradual and consistent impact as new tools become available
  - Tools can have near-term impact since ModSim precedes full development/delivery of new hardware and applications
  - Adoption of tools will depend not only on maturation of ModSim, but also readiness of programming models/environments to incorporate them.
  - ModSim can “lead” and be ready when EH programming models/environment maturity level requires ModSim cost models
  -

PRD 6.2 : Cohesive integration of modeling & simulation infrastructure with programming models/environment (continued: Jeremiah Wilke, David Donofrio)

- Metrics for success:
    - For ModSim itself: number of programming models/environments leveraging ModSim tools (and conversely number of programming frameworks providing useful application models to ModSim)
    - For EH as a whole: performance improvements obtained by leveraging cost models from ModSim
    - For EH as a whole: productivity,, “time to solution” improvements: speed at which programming tools can “down-select” from factorial design space to specific architecture/configuration
- Why make this a priority?
    - Industry tools are likely to be focused on single, proprietary architecture or specific market-driven workloads. DOE must solve EH performance portability across architectures for its workloads.
    - Success in EH ultimately depends on programming environments enabling productive application development.

PRD 6.5 : ModSim for Effective Provisioning, Allocation, and Scheduling (Andrew Chien) (2 of 2)

Metrics for Success

- Time to complete at scale simulations of new hardware tech
- Accuracy and time to evaluate where an application should run in EH facilities portfolio
- Ability to estimate what applications a new accelerator matches, and how its deployment would affect workload of various systems and resources

How and when will success impact technology?

- Embedded in procurement process for evaluation, platform scheduling software (thru vendors), and used by facilities in planning, allocation

PRD 6.4 : Application and Runtime modsim for EH (Adolfy Hoisie)

Applications in an EH environment will evolve to take advantage of new algorithms, new data structures, new interfaces for acceleration devices, including technologies BML. The complexity of optimal mapping of the application workload to system resources is anticipated to grow exponentially with the heterogeneity and the constraints of the performance/power/reliability space. To achieve that, we anticipate that runtime systems, and other layers of system software will need to be guided by models for optimization of the myriad of knobs available. Practical Modsim in this arena will require both static and dynamic (i.e., runtime) methodologies and tools.

- Research challenges
    - Capturing application workloads in models
    - New modsim techniques for application level modsim, including ML
    - Modsim for predictive design of apps and algorithms -- optimization in advance of implementation
    - Dynamic model generation based on information on the fly from a monitoring system and counters for performance, power, and thermal
    - Methods of modsim that exhibit accuracy/execution time tradeoffs
- Potential research approaches and research directions: Efforts at methodology and tool level. Optimization of apps in advance of implementation is a totally new area of investigation. Similarly, dynamic modsim for runtime system introspection.
    - Power models to enable runtime introspection
- How and when will success impact technology? Success will be gradual as new methods and tools become available.
- Why make this a priority research direction? EH exhibits a large number of degrees of freedom in the way in which mapping workloads to system resources is done for optimality. Modsim is the key to achieving productive computing by many metrics. Design space exploration. Ability to optimize execution of apps on systems dynamically.

PRD 6.5 : ModSim for Effective Provisioning, Allocation, and Scheduling (Andrew Chien) (1 of 2)

EH accelerators and complex, diverse configurations across DOE facilities will produce radically different application platform effectiveness rendering pairings from ideal to ineffective. Thus, effective matching, enabled by modsim, becomes essential component of allocation in order to effectively utilize resources and maximize scientific throughput. Effective scheduling also depends on good matching and modeling of EH resource utilization, derived from modsim, to manage provisioning and application interference management.

Research challenges:

- Modeling for systems built from new hardware technologies (quantum, analogous, etc) at scale
- Application modeling without porting to novel EH hardware systems (at scale, low cost)
- Projection/interpolation for diverse accelerators, EH features; reduce the pairings to be evaluated
- Power caps, power modeling

Research approaches and research directions

- New ModSim algorithms/tools that model performance ordering, not quantitative models
- Machine learning applied to application codes, accelerators, and facilities for creating performance estimation models
- Fast DeepNeuralNetworks/ANN (HW accelerated) and other approximation techniques for online decision-making in schedulers

PRD 6.1 : Emergent behavior and validation

As we start combining individual simulation components, how do we ensure accuracy of the system as whole since components may be used in new (untested) contexts?

- Research challenges
  - Must define metrics for progress
- Potential research approaches and research directions
  - Must show how approach can be evaluated with progress metrics
- How and when will success impact technology?
  - Must answer why DOE needs to lead aside from industry

(Maybe several viewgraphs)

Potential for merging

- 2 and 4 and 5 - Jeremiah/David
  - all have emphasis on cost models and feedback between application/runtime/OS and ModSim
  - hardware/software co-simulation important
- 1 and 3 - Adolfy/Zhiling
  - the way to reduce runtime cost is through new methodologies
  - both emphasize speed/accuracy/complexity tradeoffs
- validation still on its own - Maya/Jeanine

Challenges unique (or made worse) in EH

- Many of the same challenges as today, but worse
- Speed of simulation/modeling will become more important for larger, multi-dimensional design space
- Composability: more, different components to put together
  - more error-prone
- Validation: new challenges associated with emergent behavior putting heterogeneous components together
- Scalability:

System level modsim

- Model the interactions of interconnect networks
  - From experimental/observational data sources to high performance compute resources
- Accurately characterize overheads of task dispatch, synchronization
- Model overheads for data format conversions
- Model novel compute in wire capabilities
- Timing accurate transmission and routing in the presence of in-transit compute
- Workflow level modsim
- Workload level modsim

Things we haven’t mentioned (notes to BOG Organizers)

- Machine Learning
- Reconfigurable

ModSim Current Status - Comments from Discussion

- Status / Capability 1:

PRD X.n : Short title of possible research direction

- One paragraph description (3 sentence/bullet)
- Research challenges
    - Metrics for progress
- Potential research approaches and research directions
- How and when will success impact technology?

(Maybe several viewgraphs)

PRD 3.X

Potential approaches and directions

- ML - observation, training, automated model extraction
- Hardware acceleration of models, use parallelism, etc.
- Intelligent sampling based approaches
- Use database of previous modeling results
- Combine approaches (e.g., analytical, ML, simulation) - combine strengths of each

PRD X.3 - Reduce effective cost of modsim (1000x?)

Metrics for success

- Time (human and machine) to develop node models of different complexity
- Time to develop application models
- Time to run models and obtain meaningful results